

Detection of an orthogonal alignment between parsec-scale AGN jets and their host galaxies

Received: 3 July 2023

Accepted: 7 October 2024

Published online: 14 November 2024

 Check for updates

D. Fernández Gil^{1,2}, J. A. Hodgson¹✉, B. L'Huillier¹, J. Asorey^{3,4},
C. Saulder^{5,6,7}, K. Finner⁸, M. J. Jee⁹, D. Parkinson^{7,10} & F. Combes¹¹

The relationship between galaxies and their supermassive black holes (SMBHs) is an area of active research. One way to investigate this relationship is to compare parsec-scale jets formed by SMBHs to the projected shape of their kiloparsec-scale host galaxies. Here we analyse very-long-baseline interferometry (VLBI) images of active galactic nuclei (AGNs) and optical images of their host galaxies by comparing the inner-jet position angle in VLBI-detected radio sources with the optical shapes of galaxies as measured by several large optical surveys. In total, 6,273 galaxy–AGN pairs were found. We carefully account for the systematics of the cross-matched sources and find that Dark Energy Spectroscopic Instrument Legacy Imaging Surveys (DESI LS) data are significantly less affected by them. Using DESI LS data, with which 5,853 galaxy–AGN pairs were cross-matched, we find a weak but significant alignment signal (with $P \lesssim 0.01$) between the parsec-scale AGN jet and the kiloparsec-scale projected minor axis of the optical host galaxy in sources with measured spectroscopic redshifts. Our results show that the observed source properties are connected over three orders of magnitude in scale. This finding points towards an intimate connection between a SMBH, its host galaxy and its subsequent evolution.

One might naively expect the orientation of the supermassive black hole (SMBH)–accretion disk system to be related to the optical orientation of its host galaxy. The answer to this question (one way or another) touches on many broad areas in astronomy, such as galaxy evolution, active galactic nucleus (AGN) feedback and cosmology. Very-long-baseline interferometry (VLBI) offers a unique opportunity to probe the parsec-scale to subparsec-scale jets^{1,2}, which may be a proxy for the orientation of the SMBH–accretion disk system^{3–5}.

This question has been explored by Jones et al.⁶ with 14 galaxy–VLBI jet pairs and inconclusive results. More recently this has been looked at by Mandarakas⁷ using parsec-scale VLBI and Vazquez Najar and Andernach⁸ using kiloparsec-scale radio-jet position angles (PAs). A strong correlation between the radio jet and the optical minor axis of the host galaxies was not found; however, these two results remain unpublished. Other authors have also searched for this correlation at lower resolutions with mixed results. Mackay⁹ measured the

¹Department of Physics and Astronomy, Sejong University, Seoul, Korea. ²Centro de Estudios de Física del Cosmos de Aragón, Teruel, Spain.

³Departamento de Física Teórica and Instituto de Física de Partículas y del Cosmos (IPARCOS-UCM), Universidad Complutense de Madrid, Madrid, Spain.

⁴Departamento de Física Teórica, Centro de Astropartículas y Física de Altas Energías, Universidad de Zaragoza, Zaragoza, Spain. ⁵Max Planck Institute for Extraterrestrial Physics, Garching, Germany. ⁶Universitäts-Sternwarte München, Munich, Germany. ⁷Korea Astronomy and Space Science Institute, Daejeon, Korea. ⁸IPAC, California Institute of Technology, Pasadena, CA, USA. ⁹Department of Astronomy, Yonsei University, Seoul, Korea. ¹⁰School of Mathematics and Physics, University of Queensland, Brisbane, Queensland, Australia. ¹¹LERMA, Observatoire de Paris, UPSL Research Université, CNRS, Sorbonne Université, Collège de France, Paris, France. ✉e-mail: jhodgson@sejong.ac.kr

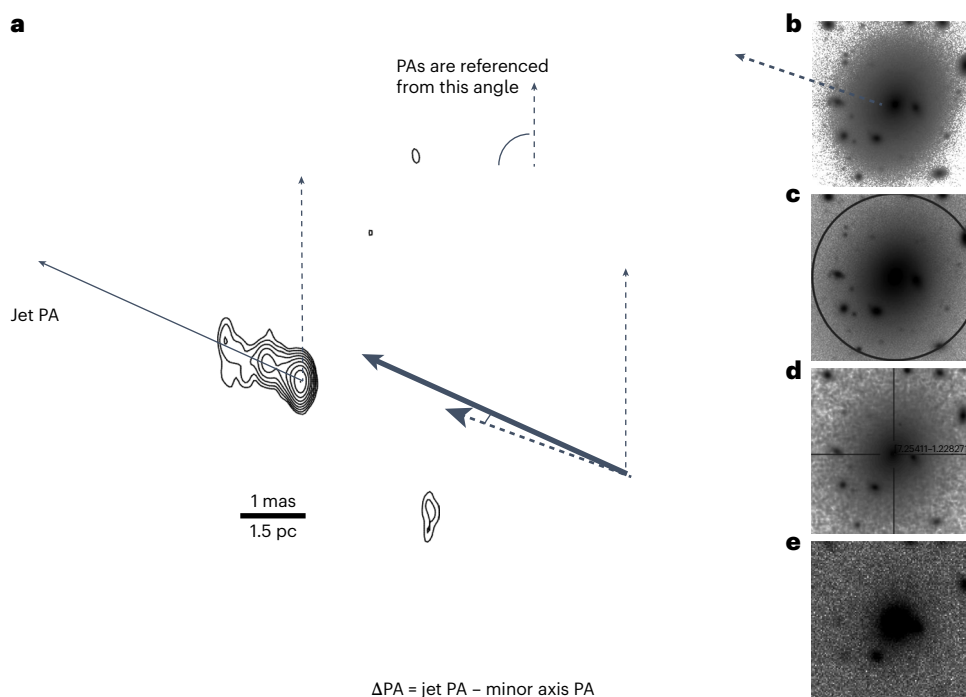


Fig. 1 | Demonstration of how the angle difference between the VLBI jet PA and the projected minor axis PA of an optical host galaxy is computed.

a, VLBI image of J0029-0113 ($z = 0.082817$) at observing frequency 8.7 GHz with an angular size of $50 \times 50 \text{ mas}^2$. **b**, Optical counterpart image from DES. **c**, Optical counterpart image from DESI LS. **d**, Optical counterpart image from SDSS.

e, Optical counterpart image from SkyMapper. All optical images have a size of $1 \times 1 \text{ arcmin}^2$. Here $1 \text{ mas} = 1.5 \text{ pc}$ ($H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.3$). The PAs are referenced north to east. We take the difference between the VLBI jet PA and the PA of the closest minor axis of the projected optical host galaxy shape.

kiloparsec-scale major axis of both the radio and optical components of elliptical galaxies, finding a strong parallel correlation between them, whereas Palimaka et al.¹⁰ used the same optical data but different radio data and found an opposite, perpendicular correlation.

Other studies looking at different, but related, topics found no alignment evidence: Kinney et al.¹¹ computed the three-dimensional angle between the radio jet and disk for Seyfert galaxies; Schmitt et al.¹² studied the three-dimensional orientation of jets relative to nuclear dust disks in 20 radio galaxies; and Greenhill et al.¹³ looked at accretion disk maser candidates, which allowed analysis of the (mis)alignment between AGNs and their surrounding galactic stellar disks.

Some works used different approaches. Lagos et al.¹⁴ simulated an AGN sample from non-AGN galaxies and found evidence that the angular momentum of material falling onto a black hole influences the AGN. Of particular relevance to our study is the work of Battye and Browne¹⁵, who compared the kiloparsec-scale 20 cm radio and optical orientations of galaxies. The optical and radio orientations were found to be perpendicular to each other in radio-quiet elliptical galaxies. In comparison, VLBI-detected sources are almost exclusively radio loud, but also hosted in elliptical galaxies^{16–19}.

However, to our knowledge, no previous study has approached this question by probing orders of magnitude closer to the SMBH–accretion disk system with VLBI while considering multiple optical surveys and carefully taking systematic errors into account.

The VLBI data were obtained from the Astroge0 VLBI FITS image database, which we use as the basis for cross-matching. Cross-matched optical data were found in the Sloan Digital Sky Survey (SDSS) Data Release 17 (DR17), the Dark Energy Spectroscopic Instrument Legacy Imaging Surveys (DESI LS) DR10, the SkyMapper Southern Sky Survey DR2 (SkyMapper), the Kilo-Degree Survey (KiDS)-1000 Gold Sample and the Dark Energy Survey (DES) DR2. For more details of the samples and surveys, see Methods. In Fig. 1, our basic method is described. We take the minor axis of the projected




Fig. 2 | Footprints of the optical surveys and the VLBI sample used in this work. Each survey is represented by a different colour, as shown in the legend.

shape of the host galaxy and the PA of the jet. The key observable is the difference between these angles (ΔPA), which lies between 0° and 90° . Care needs to be taken in transforming the PA directions into $\Delta\text{PA} \in (0^\circ, 90^\circ)$ (see ‘PA comparisons’ section). The footprints of the surveys are shown in Fig. 2. The VLBI jet PA distribution is shown in Supplementary Fig. 1, and two examples of the optical major axis PA distribution are shown: DESI LS in Supplementary Fig. 2a and SkyMapper in Supplementary Fig. 2b.

We take particular care to quantify the systematics (see Methods for a detailed description). It is unclear what is causing the systematic errors; nevertheless, we find that the DESI LS data are the least affected. For this reason, we proceed with the DESI LS data for the main conclusions. We show the results for the remaining individual optical surveys, as well as a combined dataset, in Supplementary Figs. 3–6.

Table 1 | Information summary about the optical surveys used in this work

Survey name	Median seeing value	Total cross-matches	$b > 1.3''$ ('criterion 0')	Good-case criteria	$b > 1.3''$ and good cases	$b > 1.3''$ and $\sigma_{PA} < 22.5^\circ$
SDSS	1.3" (r)	2,110	177	type= GALAXY	177	176
DESI LS	1.3"	5,853	534	type= SER, EXP or DEV	524	520
KiDS	0.7–0.9" (g)	10		B_IMAGE > 2" & CLASS_STAR < 0.5		
SkyMapper	2.6" (g)	1,337	120	b > 2" & CLASS_STAR < 0.5	56	36
DES	1.11" (g)	748	91	EXTENDED_CLASS_COADD =2,3	83	26

From left to right: name of the optical survey, median value of the seeing (the band it was measured in is in parentheses), total number of cross-matched sources, number of sources within the semi-minor axis cut ('criterion 0'), criteria for the 'good case' (see the 'Optical surveys' section in the Methods for more information), number of sources within the semi-minor axis cut and the 'good case' and number of sources within the semi-minor axis cut and the PA error cut. Text in bold denotes database parameters.

To compare VLBI jet PAs and optical PAs, it is necessary to filter out unresolved sources. To do this, we first removed any cross-matched source with a semi-minor axis of < 1.3 arcsec ('criterion 0') due to the mean seeing limits of DESI LS (Table 1). We then performed several complementary analyses. In the first case, we cut out all sources where the error in either the optical PA ($\sigma_{PA,opt}$) or the VLBI PA ($\sigma_{PA,VLBI}$) was larger than $\pm 22.5^\circ$. Any PA measurements with errors larger than this were too unreliable for our analysis. In the second case, we used various criteria for each survey (as described in Methods) to determine if an optical PA belonged to a 'good case'. Not all surveys contain the same morphological and resolution information. The criteria to consider galaxies a 'good case' need to be adapted for each individual survey, but we attempted to make them as consistent as possible. Even though the 'good case' cut could be not as homogeneous as desired, as we proceeded only with the DESI LS cross-matched sources for the main analysis, the impact of the sample selection diversity is negligible. We were motivated to perform the 'good case' and minor axis cuts to try to differentiate between sources that were resolved but mostly circular from sources that were simply unresolved. We refer to the previous three cuts as 'quality cuts'. We can also speculate that any correlation may have a redshift dependency, as closer sources would be better resolved on average. For this reason, we also make cuts with (1) all sources with spectroscopic redshifts, (2) sources with spectroscopic redshifts of < 0.5 and (3) sources with spectroscopic redshifts of < 0.1 .

In Fig. 3 we show histograms of the ΔPA distribution for two and five bins of the DESI LS cross-matches for the different cuts that were described in the previous paragraph. The left column has no 'quality cuts', the centre column contains only sources with a VLBI or optical PA error σ_{PA} below the cut and the right column contains only the 'good case' sources. In addition, each row applies a cut in redshift: the top row has no cuts, the second row contains only sources that have a spectroscopic redshift and the two lower rows contain only sources with a spectroscopic redshift of < 0.5 and < 0.1 , respectively. Each subfigure contains two histograms, one with five bins and one with two bins. Each histogram includes a dark band, which represents the expected null signal obtained by reshuffling the sample of the jet PA 1,000 times. The band is centred around the average height of each bin after the reshufflings, and the (vertical) width is the standard deviation. The histogram shows the distribution of ΔPA and the error bars of the bins are obtained from Monte Carlo simulations. The larger the deviation from the dark band, the stronger the alignment signal (that is, the P value is lower).

The two-bin histogram is justified since it tells us immediately if there are more pairs with $\Delta PA < 22.5^\circ$, that is, if we have a preference for an alignment or an anti-alignment. We investigated the effect of binning by generating results with four, five, six, eight and ten bins. We show the ten-bin histogram Supplementary Fig. 7. Changing the binning sometimes increases and sometimes decreases the P value, but does not change the overall picture. We ultimately chose to include the five-bin histogram in the main figure (Fig. 3) as it allows for a finer

analysis while keeping a statistically significant number of pairs per bin. We also performed a Kolmogorov–Smirnov (KS) test and include the P values, as it does not depend on binning. The KS test is a non-parametric test that determines the goodness of fit between two different distributions. The standard test, which is the one we use, takes a given set of data and compares it against a uniform distribution.

We further investigated the effect of magnitudes. In general, cutting by magnitude reduces the number of sources, especially in the top two rows. We demonstrate this in Methods, where we apply a strict magnitude cut in the z band of < 14 (shown in Supplementary Fig. 8) and compare the histogram with Fig. 3. As before, applying these cuts affects the number of sources, but the results remain consistent.

Our results show a weak but significant alignment ($P \leq 0.03$) in both the two- and five-bin cases for both 'quality cut' sources and 'any' sources with a measured spectroscopic redshift. No significant signal is seen when including sources without spectroscopic redshifts. The significance is highest ($P \leq 0.007$) for all sources with measured redshifts and for sources with $z < 0.5$.

In the 'good cases' cut with redshift $z < 1$ we find a median ΔPA of $40^{+34}_{-27}^\circ$ with a 68% error. We also investigated if the median evolves with redshift (Supplementary Fig. 9), but found no clear trend. Overall, we found good evidence for a weak alignment between the parsec-scale AGN jet PA and the kiloparsec-scale minor axis of the projected shape of the host galaxies for sources with a well-measured optical shape and spectroscopic redshift.

We also visualized the correlation between both angles with a scatter plot (Supplementary Fig. 10), but this does not provide a clear visible result due to the weak nature of the alignment. To help visualization, we use a colour map and add a histogram with the counts per bin to the ΔPA scatter plot in Supplementary Fig. 11.

The original assertion was that we might expect a connection between the SMBH and the general properties of the galaxy. This is particularly true in VLBI sources that are mostly hosted in radio-loud elliptical galaxies^{16,18,19}. In spiral galaxies, which are mostly star-forming galaxies, the radio emission is often due to synchrotron emission from supernovae, which traces the star-forming disk. This could explain why Battye and Browne¹⁵ reported a correlation between the optical major axis of the galaxy and the major axis of the radio images, as taken from the Faint Images of the Radio Sky at Twenty-Centimeters (FIRST) survey. Even in the case of nearby jetted spirals, an alignment of the jet with the large-scale optical disk is not expected¹¹.

A critical component to interpreting the result is the physical meaning of the projected minor axis of the host galaxy. Projection effects could be complex, but they cannot create such an alignment alone. Simulations suggest that there could be dynamical reasons why radio-loud elliptical galaxies would show a better alignment between their SMBH–accretion disk systems and their large-scale optical morphology than spiral galaxies. Spiral galaxies are mainly formed through secular evolution and gas accretion from cosmic filaments, whereas

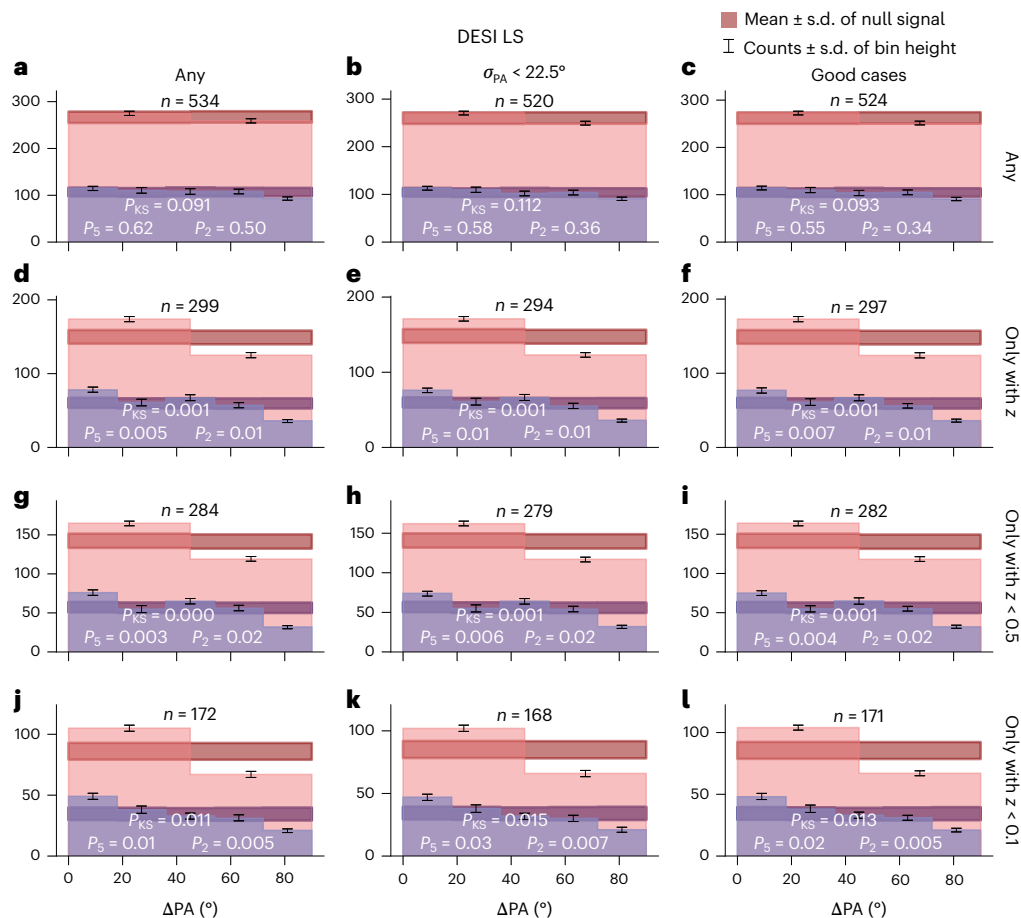


Fig. 3 | Histograms of all VLBI sources with a cross-match in DESI LS and a semi-minor axis of $>1.3''$ (the ‘criterion 0’). Δ PA indicates the angle between the VLBI jet and the optical minor axis of its host galaxy. The dark bands show the expected null signal (no alignment), which is the result of our careful analysis of the systematic errors. The stronger the alignment signal, the stronger the alignment signal. P_5 and P_2 are the P values for five bins and two bins, respectively, in the Δ PA space. P_{KS} is the P value from a KS test comparing the data against a uniform distribution, independently of the number of bins. Each panel corresponds to a specific selection of sources, showing the number of sources on top. Each panel corresponds to a specific selection of sources based on several cuts. The number of sources inside each panel is shown on top. **a**, All ‘criterion 0’

sources. **b**, All ‘criterion 0’ sources whose jet and optical PA error is $<22.5^\circ$. **c**, All ‘criterion 0’ sources within the ‘good case’. **d**, All ‘criterion 0’ sources with a measured spectroscopic redshift. **e**, All ‘criterion 0’ sources whose jet and optical PA error is $<22.5^\circ$ and with a measured spectroscopic redshift. **f**, All ‘criterion 0’ sources within the ‘good case’ and with a measured spectroscopic redshift. **g**, All ‘criterion 0’ sources with $z < 0.5$. **h**, All ‘criterion 0’ sources whose jet and optical PA error is $<22.5^\circ$ and with $z < 0.5$. **i**, All ‘criterion 0’ sources within the ‘good case’ and with $z < 0.5$. **j**, All ‘criterion 0’ sources with $z < 0.1$. **k**, All ‘criterion 0’ sources whose jet and optical PA error is $<22.5^\circ$ and with $z < 0.1$. **l**, All ‘criterion 0’ sources within the ‘good case’ and with $z < 0.1$.

elliptical galaxies are often formed by major mergers²⁰. These different formation scenarios have alignment consequences with cosmic filaments²¹. Simulations also suggest that the remnant of a major merger would be a spheroidal object elongated along the arrival direction of the swallowed companion, with a small disk forming from the tidal debris, slowly rotating around one of the minor axes of the remnant. The gas from this disk will be accelerated from the SMBH while conserving its angular momentum, resulting in a jet that is perpendicular to the accretion disk. In spiral galaxies, both cosmic gas accretion and minor mergers occur isotropically, which results in a final random accretion direction onto the SMBH. Small radio jets in spirals often encounter disk material and are bent several times, so that the outflow eventually escapes from the disk material in a perpendicular direction at large scales.

We can use the aforementioned simulations to better understand the physical meaning of the signal. We used the Evolution and Assembly of GaLaxies and their Environments (EAGLE) simulation^{22,23} and explored four scenarios: (1) no modifications to the simulation data with a close three-dimensional relationship between the stellar component spin of the galaxy and the main plane of the galaxy (see Methods

for more details), (2) randomizing the three-dimensional relationship, (3) adding a Gaussian scatter and (4) adding a uniform scatter. For scenarios 3 and 4, the Gaussian and uniform scatters were added to the original close relationship of 33° , as motivated by the 68% errors on the median of the Δ PA distribution in the DESI LS ‘good cases’ with $z < 1$. In the first scenario (Fig. 4a) we get a very strong alignment signal. In the second scenario (Fig. 4b) we recover a uniform signal, as expected, and in the third scenario (Fig. 4c) we interestingly recover a distribution that looks qualitatively similar to our observed distributions in the DESI LS ‘quality cuts’ with low redshift, with a similar median and 68% error in Δ PA. The fourth scenario (Fig. 4d) shows a slightly stronger signal.

While the simulations could be suggestive, we should be careful not to over-interpret as there are many reasons why such an alignment could weaken. The direction of jets may depend on the radio observing frequency and can vary on timescales of the order of AGN duty cycle, that is 10 Myr (ref. 24), much smaller than the timescale of optical morphology variation for galaxies (several billion years). The jet direction depends on the angular momentum of the accreted material, and gas can come from a gas-rich late minor merger. Similarly, because mergers are complex dynamical processes, the angular momentum connection

between the SMBH and the kiloparsec-scale structure could be scrambled, for example²⁵. One possibility that may result in a weakened minor axis alignment, and which has been seen in some cases, for example²⁶, is that the optical jet or dust dominates the morphology of the host galaxy, which explained the previously observed offset in VLBI and Gaia sky positions^{27,28}. In this case, the optical shape would correspond to the jet rather than that of the galaxy, and the inferred galaxy PA would be the direction of the jet. This could yield an alignment with the major axis. In the case of dust emission, it can also affect positional offsets between VLBI and Gaia, which could also affect the alignment signal reported here. All else being equal, complicated effects would be expected to weaken any observed signal.

Simulations also show that misalignments arise naturally from the angular momentum transfer; therefore, only a weak connection would be predicted, and even then only for maximally spinning black holes²⁹.

A connection between the morphologies over such large scales could support the hypothesis of co-evolution of the black hole and the host galaxy²⁰; however, the interpretation of the result in this context remains unclear.

If the parsec-scale jets are indeed a proxy for the SMBH–accretion disk system, then this points to connections over an even larger range of scales. Observationally, the link between the SMBH and the inner jet sometimes seems clear: Gómez et al.³⁰ probed the innermost jet of BL Lacertae with the RadioAstron space VLBI antenna, Boccardi et al.³¹ probed the two-sided jet in Cygnus A with the Global mm-VLBI Array (GMVA) and Event Horizon Telescope observations of both M87 and Cent A^{32–34}. However, other studies were not so clear: Hodgson et al.³⁵ used the GMVA and Gómez et al.³⁶ used RadioAstron to find parsec-scale jet bends in OJ 287, Giovannini et al.³⁷ found large cross jet opening angles in 3C 84 using RadioAstron, which was confirmed using the GMVA by Oh et al.³⁸ and Paraschos et al.³⁹. Similarly, Kim et al.⁴⁰ used the Event Horizon Telescope to find subparsec-scale jet bends in 3C 279. Perhaps crucially, however, further downstream (parsec-scale) the jets seem to follow the overall jet direction. The connection of the SMBH to the accretion disk system itself is also unclear, for example⁴¹, and will require further study. Nevertheless, the correlation could suggest that the momentum of the SMBH–accretion disk system is influencing or being influenced by the host galaxy and/or that the correlation is more influenced by the interplay between the jet itself and the host galaxy (that is, AGN feedback).

An important finding is that having a well-measured shape and spectroscopic redshift introduces the alignment signal. We see in the last row of Fig. 3 that the signal is most significant, but the improvement is not so large. This is largely to be expected as the number of sources in that row also reduces. Furthermore, the AGN density peaks around $z \sim 1-2$; therefore, there are not many radio-loud AGNs at $z < 0.1$. Another reason is that at low z we observe smaller volumes. Having a measured redshift introducing the signal suggests a strong resolution effect as sources with $z > 0.5$ rarely have well-measured shapes. This is reflected in the relatively small change in source counts between any redshifts and $z < 0.5$.

Our examination of the ‘quality cuts’ seems to indicate a scarcity of genuinely spherical galaxies associated with a radio-loud jet. This implies that nearly all radio-loud elliptical galaxies possess a non-spherical form, potentially linked to the SMBH and accretion disk system. If this linkage is also discovered in spiral galaxies and radio-quiet AGNs, it could point towards a more universal correlation. However, it is difficult to probe the orientation of the SMBH–accretion disk system without jets.

Conclusion

We find good evidence that the parsec-scale VLBI jet is connected to the projected kiloparsec-scale optical host galaxy morphology. The direction of the parsec-scale jet (which may be a proxy for the SMBH–accretion disk system) appears to be oriented perpendicularly to the

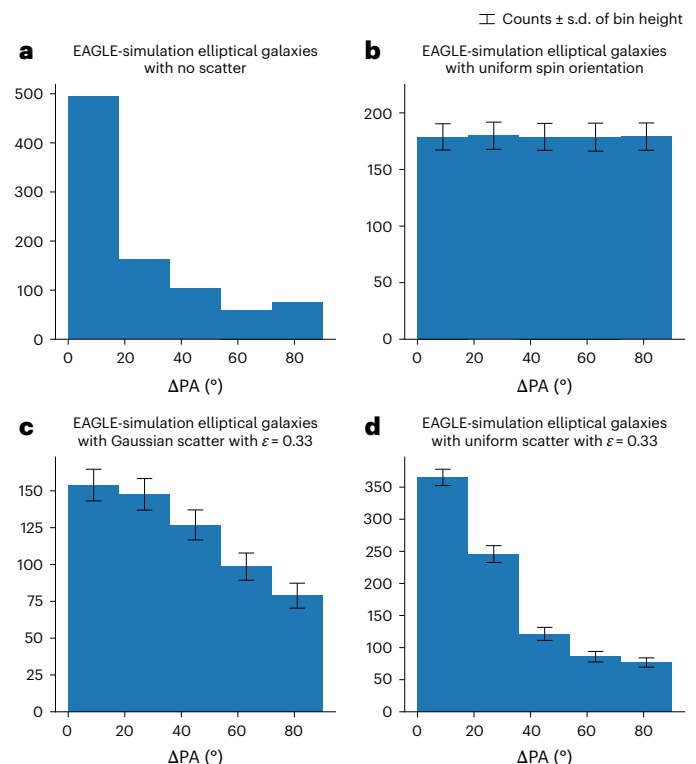


Fig. 4 | ΔPA between the projected spin axis of the stellar component and the semi-minor axis of the projected ellipse of the elliptical galaxies in the EAGLE simulation. ϵ indicates the amplitude fraction of the offset we introduce to the three-dimensional spins. The histogram is the result of averaging 1,000 simulations, each time sampling the spin tilt offset errors randomly from different distributions, as explained in the text above and in Methods. The error bars are the standard deviations. **a**, No offset. **b**, Uniform spin orientations. **c**, Gaussian scatter with $\epsilon = 0.33$. **d**, Uniform scatter with $\epsilon = 0.33$.

projected optical shape of the host galaxy with a P value of < 0.05 . The interpretation of these results is complicated but may have far-reaching implications. Ultimately, these results will need to be followed up with higher source counts and higher-resolution optical observations (for example Hubble Space Telescope, James Webb Space Telescope) and with radio-quiet VLBI observations in order to determine how universal such a connection may be. In the near future, large surveys will be performed, for example those using the Vera C. Rubin Observatory and Square Kilometer Array VLBI, which have the potential to vastly increase our sample sizes.

Methods

Observations

VLBI jet images and the Astrogate catalogue. VLBI is a technique that allows for radio antennae, which can be separated by many thousands of kilometres (or even into space), to be combined to create ultra-high-angular-resolution images. This angular resolution is set by the ratio λ/D , where λ is the observing wavelength and D is the distance between the most spatially separated antennae. VLBI has been used to image black hole shadows³² and to measure cosmic distances⁴².

The Astrogate VLBI image database contains over 112,000 images from more than 17,000 sources, which have been collected from various VLBI surveys, for example^{43,44}. The observations are performed at observing frequencies ranging from 1.4 to 86 GHz. Most observations are performed at 2 GHz, 5 GHz, 8 GHz or 15 GHz. Because these observations use globally located antennae, angular resolutions of sub-milliarcseconds are regularly achieved. These angular resolutions often correspond to subparsec scales even at relatively high redshifts

($z > 1$). In Fig. 1, an example VLBI source at parsec scale is displayed. This scale allows us to probe the central black hole–accretion disk system of the galaxy. In contrast, the optical images at the corners are taken from the various optical surveys (described in more detail in the next section) and show the host galaxy in vastly larger (kiloparsec) scales.

Typically, the sources detected in the VLBI observations are the bright, compact cores of radio AGN jets. These sources frequently exhibit a core-jet structure where the core is the bright component most upstream in the jet and the jet emission is fainter, extended and emanates from the core.

In cases where there is a clear core-jet structure, a jet PA, that is, the direction of the jet emission relative to the core, can be reliably determined. This analysis was performed by Plavin et al.⁴⁵, who found reliable PAs for 9,220 sources. The jet PAs were measured by fitting two Gaussian components directly in the u – v plane using a Bayesian nested sampling algorithm, which also provides formal errors. They then compared these formal errors against the intrinsic scatter over multiple epochs and/or frequencies in order to estimate the true errors.

These sources form the basis of our VLBI dataset. Redshifts were obtained from the Optical Characteristics of Astrometric Radio Sources (OCARS)⁴⁶ catalogue. The catalogue differentiates between spectroscopic, photometric and generally unreliable redshift measurements. When redshifts were used, we used only spectroscopic redshifts. These were found for roughly 75% of sources with a redshift measurement, and 43% of all sources.

Optical surveys. For the orientation of the host galaxies, we used photometric images from five optical surveys. These surveys are listed in Table 1: SDSS, DESI LS, KiDS, SkyMapper and DES. In all cases, optical shapes are obtained from catalogues provided by these surveys. Figure 2 shows the footprint of each survey, and details are given below.

Sloan Digital Sky Survey DR17. The SDSS is an optical survey with imaging and spectroscopic information covering roughly 14,500 square degrees of the sky in both hemispheres^{47,48}. We use the latest data release, DR17 (ref. 49), which belongs to phase IV of the survey. SDSS used the Sloan Foundation 2.5 m Telescope at Apache Point Observatory in New Mexico, USA. The imaging data used here were taken by the SDSS imaging camera⁵⁰ taking data in the u , g , r , i and z bands.

We accessed the galaxy photometric information through the PhotoObjAll table, which provides, among others things, the magnitude, the PA of the major axis and the minor–major axis ratio in five different bands: u , g , r , i and z . SDSS fits two models to each object in each band: a pure de Vaucouleurs (DeV) profile and a pure exponential (exp) profile. Each model has an arbitrary axis ratio and PA. The models are convolved using a double Gaussian fit to the point spread function (PSF) and the likelihood associated with the χ^2 fit is computed. Although for large objects it is possible, and even desirable, to fit more complicated models (for example, bulge plus disk), the computational expense to compute them is not justified for the majority of the detected objects⁴⁹. They provide the PA (Φ) and axis ratio (AB) parameters for each model, plus the modelMag parameter, which is the magnitude corresponding to the best DeV/Exp magnitude fits in the r band. For our purposes, we use the best of these two models only in the r band. So, we would access the PA of a specific source in the r band using the `deVPhi_r` parameter if the best-fit model is DeV or the `exPhi_r` parameter if it is Exp. To make sure that the choice of band did not strongly affect the results, we compared the optical PA between all the different bands; they are mostly correlated, with some having an offset of 180° due to the galaxy symmetry. The errors on the major axis PAs are not provided since DR8 because systematic errors dominate, over the photon statistical errors. These systematic errors are of the order of 15° (ref. 51). We adopted this value for all SDSS sources analysed in this study. Finally, we also made use of the morphological classification type (3 for resolved or GALAXY, 6 for unresolved or STAR)

of the object, which was computed by comparing the magnitudes of the PSF and the fitted model.

SkyMapper Southern Sky Survey DR2. The SkyMapper Southern Sky Survey is a wide-field optical survey that covers >21,000 square degrees using the SkyMapper 1.35 m telescope at Siding Springs Observatory in Australia. Using its own wide-angle 268 megapixel camera, the survey focused on creating a full survey of the southern sky. In particular, DR2 contains up to 500 million objects between AB magnitudes 8 and 22 in the u , v , g , r , i and z bands^{52,53}. We used the DR2 photometry table to obtain the shape parameters. Petrosian models were used to fit the parameters to the galaxies. Once the Petrosian radius was computed, its PA and axis ratio were obtained. The major axis PA was provided in the `pa` parameter. Unlike in SDSS, the error information of the PAs is provided in the `e_pa` parameter and was included in our analysis. This catalogue does not provide a systematic classification of objects between point-like and extended, but it provides CLASS_STAR, the stellerity index from SExtractor; and the semi-minor axis `b`. CLASS_STAR defines an a posteriori probability of a detection to be a point source or an extended object. It is usually closer to zero for extended sources and close to one for point-like sources⁵⁴.

DESI Legacy Imaging Surveys DR10. The DESI LS⁵⁵ consist of three separate projects: the Dark Energy Camera Legacy Survey (DECaLS), the Beijing–Arizona Sky Survey (BASS) and the Mayall z -band Legacy Survey (MzLS). The data from BASS and MzLS, which cover a large part of the north galactic cap north of a declination of 32.375°, are the same as those from the previous iteration of the DESI LS DR9. The current data release, DR10, has been supplemented by new data observed with the Dark Energy Camera (DECam), mainly on the southern hemisphere. This includes data from the DES, the DECam Local Volume Exploration Survey and the DeROSITA Survey. The surveys conducted with the DECam in the g , r and z bands were carried out on the Blanco 4 m telescope, located at the Cerro Tololo Inter-American Observatory in Chile. The BASS was carried out using the 90Prime camera with g - and r -band filters at the Bok 2.3 m telescope, and the MOSAIC-3 camera with a z -band filter at the 4 m Mayall telescope was used for the MzLS. The Bok and Mayall telescopes are located next to each other at Kitt Peak National Observatory in Arizona, USA.

With this wealth of data, the Tractor⁵⁶ pipeline was used to obtain photometric models for all detected sources across multiple imaging bands. These galaxy model parameters were obtained directly from Tractor (Lang et al., in preparation), which fits simple elliptical models consisting of single elliptical isophotes and suitable radial brightness profiles for the various morphological types (REX, EXP, DEV and SER) to the DESI LS. Using these models, the parameters are measured by optimizing a pixelwise forward model, correlating each model image with the individual image PSFs. Using the SWEEP catalogues created by Tractor, we obtained fluxes and shape parameters like the PA and its errors. As the shape measurements were obtained jointly across all available optical bands, we selected only objects from which the catalogue had non-zero fluxes in the g , r and z bands with finite measurement uncertainties.

To each of the detected sources, Tractor fits a range of morphological models: PSF (stellar, unresolved), REX (round exponential galaxy with a variable radius), EXP (exponential), DEV (de Vaucouleurs), SER (Sérsic) and DUP (sources detected by the Gaia Space Telescope but not by DESI LS). The morphological type used for the set of parameters provided for each source is determined by the T value of the model, which also gets penalized (<https://www.legacysurvey.org/dr9/description/#morphological-classification>) for more ‘advanced models’, with PSF being the simplest, then REX, followed by EXP and DEV and finally SER as the ‘best’ model. The DEV, EXP and SER models provide shape fits for extended sources and with them: Sérsic indices n , half-light radii r_{eff} and the complex ellipticities ($e_1 + i e_2$), from which the

axis ratio $\frac{b}{a}$ and axis PA ϕ , and their errors, can be obtained using the following equations:

$$\bar{\epsilon} = \sqrt{\epsilon_1^2 + \epsilon_2^2}, \quad (1)$$

$$\frac{b}{a} = \frac{1 - \bar{\epsilon}}{1 + \bar{\epsilon}},$$

$$\phi = \frac{1}{2} \arctan \left(\frac{\epsilon_2}{\epsilon_1} \right). \quad (2)$$

Kilo-Degree Survey-1000 Gold Sample. The KiDS is an optical wide-field imaging survey designed to measure weak gravitational lensing using OmegaCAM on the European Southern Observatory's Very Large Telescope (VLT) Survey Telescope at Cerro Paranal in Chile. It has a sky coverage of 1,500 square degrees, with imaging data in five near-infrared bands, Z , Y , J , H and K , and four optical bands, u , g , r and i . The KiDS-1000 dataset contains only galaxies with reliable shape and redshift measurements, their 'gold sample'. Photometric shape parameters were provided by using their Astro-WISE image processing pipeline, which creates co-added images combining all bands, and then applying SourceExtractor to them. SourceExtractor works by carrying out isophotal measurements on the filtered, background-subtracted detection images. It uses only pixels with values above the detection threshold and fits a single ellipse into them, with the corresponding PA and axis ratio. This is a very quick process, but it is very dependent on the detection threshold itself. As mentioned before, more complex models could be fitted to some of our galaxies, but this is outside the scope of the paper, which aims to focus on a first detection of correlations. The major axis PA and its error are provided in the THETA_J2000 and ERRTHETA_J2000 parameters, and we use the stellarity index CLASS_STAR (another parameter provided by SourceExtractor) and the minor axis B_IMAGE to determine the resolution of the sources, as in SkyMapper (see 'SkyMapper Southern Sky Survey DR2' section and de Jong et al.⁵⁷). Only ten VLBI sources were found to cross-match with the KiDS data.

Dark Energy Survey DR2. The DES is a ground-based, wide-area, optical and near-infrared imaging survey that took place from 2013 until the last observing season in 2019. For this, the DES collaboration designed, built and used the DECam, a 570 megapixel camera at the 4 m Blanco telescope in Chile^{58,59}. For their cosmology goals, DES implemented two surveys: a main wide survey covering about 5,000 square degrees in the grizY photometric bands and a 30-square-degrees deep supernova survey in the griz bands. The data are processed at the National Center for Supercomputing Applications using the DES pipeline⁶⁰. For this paper, we used the DES DR2 (refs. 61,62).

Shape measurements were provided in the DR2 MAIN table, which is accessible through the DESaccess page (<https://des.ncsa.illinois.edu/desaccess/>). They obtain these parameters by applying SourceExtractor to the co-added images, which are combined from different epochs and different bands $r+i+z$ (refs. 60,61). We specifically used the THETA_J2000, the PA of the source in J2000 coordinates computed by SourceExtractor as explained above, and ERRTHETA_IMAGE, its error. The PA errors are higher than those of DESILS, even though the camera is the same. This may be because DES relies on SourceExtractor, which fits a single ellipse, whereas DESILS uses its own TRACTOR catalogue fitting more complex models. We also utilized the EXTENDED_CLASS_COADD parameter, based on the SPREAD_MODEL parameter from SourceExtractor, which compares the PSF with a slightly extended circular exponential convolved with the PSF. Stars and point-like objects have values of SPREAD_MODEL very close to zero, whereas extended sources have higher values. EXTENDED_CLASS_COADD is defined as a sum of three Boolean conditions involving SPREAD_MODEL in the i band. Every condition that is true adds one to the final value, so that

EXTENDED_CLASS_COADD = 0 defines high-confidence stars, EXTENDED_CLASS_COADD = 1 defines likely stars, EXTENDED_CLASS_COADD = 2 defines likely galaxies and EXTENDED_CLASS_COADD = 3 defines high-confidence galaxies⁶³.

Underlying distribution of the galaxy and jet PAs

Before we started using the VLBI and optical PAs, we checked for systematics in their distribution. While some surveys show a uniform distribution of PAs, others show non-uniform distributions, which may be caused by unaccounted-for systematics. We show the distribution of the VLBI jet PAs in Supplementary Fig. 1, and as an example we show the two most extreme cases of the optical PAs in Supplementary Figs. 2a and 1b: DESILS with a very uniform distribution, and SkyMapper with a completely skewed distribution.

There are slightly more VLBI jets pointing towards the vertical direction (with a PA of either 0° or 180°). Mandarakas et al.⁶⁴ show that this effect may be due to a north-south elongation effect from VLBI observations with an elliptical beam, especially affecting sources that are located in the southern region of the sky (dec. < -30°). However, the PA measurements used in this paper were measured directly from the interferometric visibilities; therefore, the PA measurements should not be affected in this way⁴⁵. The optical counterparts of SkyMapper present this strange preference for the same direction as well (an optical PA of 0°), whereas this effect was not seen in a random sample of galaxies from the same survey.

We describe how to account for this effect when quantifying the alignment signal below.

Errors

Because VLBI source positions are very precise (milliarcsecond precision or better), we used their positions as the basis for cross-matching. We cross-matched optical positions within a tolerance of 1 arcsec (which is approximately equivalent to the maximum seeing limit of the optical surveys). This yielded a total of 6,273 sources that had optical cross-matches in at least one survey.

We obtained the errors in Δ PA by adding the error in the VLBI jet and in the optical axis PA in quadrature, assuming they are uncorrelated:

$$\sigma_{\Delta\text{PA}} = \sqrt{\sigma_{\text{VLBI PA}}^2 + \sigma_{\text{Optical PA}}^2} \quad (3)$$

To obtain the error bars of the bins in the Δ PA histograms, by assuming that the errors of the jet and galaxy PAs were Gaussian, we simulated their distributions by redrawing the PA in each survey 10,000 times. We now call these errors $\sigma_{\Delta\text{PA}}^{\text{MC}}$ for Monte Carlo.

Cuts

Semi-minor axis cut. We imposed a semi-minor axis b cut, keeping only sources with $b > 1.3''$ (criterion 0) based on the mean DESILS seeing value (Table 1) and applied it to all surveys. In general, this cut removes many sources from the 'all cases' column, but also removes some sources in the 'quality cuts', particularly in the 'only with z' ' and $z < 0.5$ rows; however, it does not affect the final results.

'Good case' cuts. In the main text, the 'quality' and redshift cuts were described. In this section we provide some more relevant details.

To define the 'good case' cuts, we use various criteria from each optical survey to determine if an optical PA belongs to it, as listed in Table 1 and defined as follows:

- SDSS: class = 3 sources (where 3 are extended and 6 are point-like)
- DESILS: type = EXP, DEV and SER sources with observations in g , r and z bands and non-zero flux in all of the bands (where EXP is an exponential, DEV is a de Vaucouleurs and SER is a Sérsic model, while PSF is a point-like object)

- SkyMapper: minor axis $b > 2$ arcsec and stellar index $\text{CLASS_STAR} < 0.5$ (where $\text{CLASS_STAR} = 0$ corresponds to ‘not stars’ and $\text{CLASS_STAR} = 1$ corresponds to ‘stars’)
- KiDS: minor axis $B_IMAGE > 2$ arcsec and stellar index $\text{CLASS_STAR} < 0.5$ (similar to SkyMapper)
- DES: $\text{EXTENDED_CLASS_COADD} \in \{2, 3\}$ (where 2 corresponds to ‘mostly galaxies’ and 3 to ‘high-confidence galaxies’)

In the SDSS data, all errors are set to 15° ; this means that there may be many unreliable PA measurements that passed the PA error cut. This is an additional motivation to use the ‘good case’ cut. Sources that are cross-matched with optical counterparts and lack a spectroscopic redshift measurement are uploaded in an online repository.

Magnitude cuts. We performed a magnitude cut for our DESI LS analysis. We performed cuts for $z < 23$, 20, 17 and 14 and we include the most severe ($z < 14$) in Supplementary Fig. 8. The number of sources in the first row decreases significantly. The second and third rows are identical in all cases. This tells us that all sources with redshifts and below magnitude $z < 14$ have well-measured shapes. The final row is only marginally different. However, overall, the significance of the alignment signal remains consistent with that of Fig. 3.

PA comparisons

It would seem that the best way to visualize the alignment between VLBI jets and a galaxy’s minor axis is to check their correlation in a scatter plot. However, this method poses several problems.

The main issue is that this plot would not take the circular nature of the angles into account. The visualization of the angular distance between two directions in a linear plot can get tricky if one imposes the wrong constraints to the range of each angle. We especially have to consider that the semi-minor axis has a 180° rotational symmetry, while the jet has one specific direction. The VLBI jet PA is defined within a $(-180^\circ, 180^\circ)$ range, to cover the full range of directions, while the galactic minor axis is within a $(-90^\circ, 90^\circ)$ range, which is enough due to its symmetry. If we made a scatter plot with these default ranges, there would be a misleading representation of the true alignment. If, for example, a jet has a $\text{PA} = 87^\circ$ and the minor axis has a $\text{PA} = -87^\circ$, the point would fall far away from the diagonal (in one of the corners in the scatter plot), while the real angular difference between the two is $\Delta\text{PA} = 6^\circ$, very close to aligned. To handle this, we take the PA of the semi-minor axis and add or subtract 180° when needed until its numerical value is always within a $(0^\circ, 90^\circ)$ separation from the numerical value of the jet PA. In the previous example, the semi-minor axis PA would become 93° . This turns the semi-minor axis range to $(-270^\circ, 270^\circ)$, and the points will always fall within a $+90^\circ$ band of the diagonal. Supplementary Fig. 10 shows this computation for two examples: one with all sources with semi-minor axis $b > 1.3''$ where there is no signal, and one with the good case with $b > 1.3''$ and $z < 0.5$. The important point is that we can directly recover the ΔPA histograms (on the right) from the scatter plots (on the left) by taking $\Delta\text{PA} = y - x$. We have illustrated that by putting the histograms to the right of the scatter plots.

However, even here, we see that the alignment signal is not strong enough for the correlation to be easily visible. To aid with the visualization, Supplementary Fig. 11 shows a coloured two-dimensional histogram, to help show the excess (or lack of excess) of sources in a given bin.

Quantifying the alignment signal

The underlying distributions of PAs of some of the surveys may not be uniform. This lack of uniformity might influence the results and even create a false alignment signal. We wanted to make sure that the signal we found was not affected by this. In order to do that, we studied how far the observed signal was from what we would see if the PAs were

taken at random from their underlying distributions. This approach is commonly used to detect alignment signals^{65,66}.

This analysis is applicable to both the two- and five-bin cases. In order to do this, we shuffled the underlying PA distributions of both the jets and the optical minor axes $n_{\text{RS}} = 1,000$ times. For each reshuffling (RS) j , we recalculate the histogram of the reshuffled PA differences, $\Delta\text{PA}^{\text{RS},j}$. The height of each histogram bin after doing the reshuffling is $n_{i,j}$, where i denotes the bin and j the reshuffling. For each bin i we calculate the mean ($\bar{n}_{\text{RS},i}$) and the standard deviation ($\sigma_{\text{RS},i}$) of the number counts from all the reshufflings, defined as

$$\bar{n}_{\text{RS},i} = \frac{1}{n_{\text{RS}}} \sum_{j=1}^{n_{\text{RS}}} n_{i,j} \quad (4)$$

$$\sigma_{\text{RS},i}^2 = \frac{1}{n_{\text{RS}}} \sum_{j=1}^{n_{\text{RS}}} (n_{i,j} - \bar{n}_{\text{RS},i})^2, \quad (5)$$

where $\bar{n}_{\text{RS},i}$ is the centre of the dark bands and $\sigma_{\text{RS},i}$ is the (vertical) width of those bands present in each bin of the histograms in Fig. 3 and Supplementary Figs. 3–8. These bands, as explained, reflect the distribution of the underlying data. Therefore, if there are systematics in the data, this would be reflected as a deviation from a flat band.

Next, we defined a statistic (Q) to quantify the significance of any possible signal away from a random occurrence. We computed Q by using the bin heights of the original, unshuffled histograms ($n_{\text{data},i}$) and comparing them with the mean and standard deviation of the reshuffled data (the dark bands) as defined before. This is called Q_{data} :

$$Q_{\text{data}} = \frac{1}{n_{\text{bins}}} \sum_{i=1}^{n_{\text{bins}}} \left[\frac{n_{\text{data},i} - \bar{n}_{\text{RS},i}}{\sigma_{\text{RS},i}} \right]^2, \quad (6)$$

Q_{data} is thus a quantification of the deviation from the expected random signal as computed by reshuffling. A low Q_{data} corresponds to consistency between the measured and the randomized signals, whereas a true alignment signal would yield a large Q_{data} .

In order to determine the statistical significance, we calculated Q for each of the 1,000 reshufflings j as:

$$Q_j = \frac{1}{n_{\text{bins}}} \sum_{i=1}^{n_{\text{bins}}} \left[\frac{n_{j,i} - \bar{n}_{\text{RS},i}}{\sigma_{\text{RS},i}} \right]^2. \quad (7)$$

where $n_{j,i}$ is the height of bin i in each reshuffling j .

We now calculate the P value, that is, the probability that any given Q_j is greater than Q_{data} under the null hypothesis H_0 of non-alignment $P(Q_j > Q_{\text{data}} | H_0)$. This shows how far the Q_{data} lies from the distribution of the 1,000 Q_j from all the reshufflings.

Effect of binning

We investigate the behaviour of the results using different numbers of bins (two, four, five, six, eight and ten) for the histograms. We found that, even for higher bin numbers, the P values remain consistent, sometimes increasing and sometimes decreasing, but not changing the overall picture. We present the most severe case of ten bins (and the two-bin case as a reference) in Supplementary Fig. 7.

Redshift evolution

We show the median ΔPA versus redshift of the DESI ‘good cases’ with a redshift $z < 1$ in Supplementary Fig. 9, generating redshift bins and taking the median ΔPA of the sources in each bin. We overplotted the number of sources inside each bin. We found that, in general, both the mean and the median ΔPA were below 45° , particularly in redshift bins with high source counts. We found no clear trend with redshift. Overall, the median and 68% errors are 40^{+34}_{-27} .

EAGLE simulation projection effects

To investigate how projection effects could affect the observed two-dimensional ΔPA distribution, we used data from the 100 Mpc box of the EAGLE simulation^{22,23}. This cosmological hydrodynamical simulation provides catalogues⁶⁷ for a wide range of properties of galaxies. This includes the morphology and kinematics of galaxies⁶⁸, profile fits to artificial images⁶⁹ created from the simulated galaxies that provide the PAs along the z axis of the simulation and the colours and luminosities⁷⁰ of said galaxies. Furthermore, the EAGLE simulation provides the three-dimensional spin vectors of several components of the galaxies (stars and cold and hot gas).

The three-dimensional stellar spin vector is assumed to be close to perpendicular to the main plane of the galaxy. To confirm this, we split the sample into ellipticals and spirals using a cut in the $u-r$ colour ($u-r > 2.2$ for ellipticals, $u-r < 2.2$ for spirals). We saw no significant differences in the results from both samples; nevertheless, we used the elliptical sample. To emphasize, we are not implying a physical connection between the stellar spin and the jet direction, we are just investigating projection effects. We then projected the stellar spin onto the same plane as the projected PA of the galaxies and compute the two-dimensional projected ΔPA between the minor axis and the spin, producing histograms in a style similar to that of Fig. 3.

We studied four different scenarios, as mentioned in the main text: (1) simulation data with no additional errors, (2) fully randomized stellar spin directions and (3) and (4) added scatter to the spin directions. We provide more details below.

To account for differences between a ‘pure’ and a more realistic connection, we introduced a scatter in the three-dimensional spins before projecting them. We drew angles from both a Gaussian (3) and a uniform (4) distribution for θ , where $\theta \in N(0^\circ, \epsilon 90^\circ)$ and $\theta \in U(0^\circ, \epsilon 90^\circ)$, and a uniform distribution for ϕ in both cases, where $\phi \in U(0^\circ, 360^\circ)$. Here, $\epsilon \in (0, 1)$ represents the fractional scatter. We then tilt the spin by an amplitude angle θ and a random orientation ϕ . Our dataset consists of approximately 3,000 galaxies from the simulation. To generate a ΔPA histogram, we repeated the process 1,000 times, each time sampling the angles randomly from the distributions. Finally, we calculated the average and standard deviation for each bin in the ΔPA space and plotted the final histogram. We show the resulting histograms in Fig. 4.

Combined histogram

If the VLBI source was cross-matched to multiple optical surveys, we computed the weighted average of the angle difference ($\overline{\Delta PA}$) and its error ($\sigma_{\overline{\Delta PA}}$) in quadrature:

$$\overline{\Delta PA} = \sum_i w_i \Delta PA_i, \quad (8)$$

$$\sigma_{\overline{\Delta PA}} = \sqrt{\frac{1}{\sum_i \sigma_{\Delta PA_i}^2}}, \quad (9)$$

where i refers to the optical survey and w_i , ΔPA_i and $\sigma_{\Delta PA_i}$ are the corresponding weight, angle difference and its error. The weights are defined as:

$$w_i = \frac{1/\sigma_{\Delta PA_i}^2}{\sum_i 1/\sigma_{\Delta PA_i}^2}. \quad (10)$$

In Supplementary Fig. 3, we present the distributions of the weighted angle between the VLBI jet PA and the minor axis of the apparent orientation of the host galaxy ($\overline{\Delta PA}$). The histogram is divided in the same way as described in the main text (for example Fig. 3). In general, the combined histogram follows the DESI LS histogram in Fig. 3, as described in the main text. As before, we can see several trends in the data. In the left column, containing all sources, we do not see any

alignment overall ($0.3 < P < 0.5$), but, similarly to the DESI LS histogram (Fig. 3), we see a weak alignment for $z < 0.1$. When we look at the ‘good’ cases, we begin to see a trend towards $\overline{\Delta PA} = 0^\circ$ (that is, the inner-jet PA aligned with the minor axis of the projected optical orientation). This trend becomes slightly more significant if the sources have spectroscopic redshifts ($0.1 < P < 0.3$), and is especially strong for sources with spectroscopic $z < 0.1$ ($0.01 < P < 0.1$).

Survey-by-survey analysis

We also performed the analysis on a survey-by-survey basis; the results are shown in Supplementary Figs. 4–6. We do not include the individual analysis of the KiDS catalogue here due to the very low number of cross-matches.

Data availability

The data and codes for reproducing the results of this work are available via GitHub at <https://github.com/davfer12/Detection-of-an-orthogonal-alignment-between-parsec-scale-AGN-jets-and-their-host-galaxies>. In the repository, we provide data for the cross-matched sources, but we do not include the whole catalogues as they are publicly available from their own repositories.

References

- Weaver, Z. R. et al. Kinematics of parsec-scale jets of gamma-ray blazars at 43 GHz during 10 yr of the VLBA-BU-BLAZAR program. *Astrophys. J. Suppl. Ser.* **260**, 12 (2022).
- Pushkarev, A. B. et al. MOJAVE - XX. Persistent linear polarization structure in parsec-scale AGN jets. *Mon. Not. R. Astron. Soc.* **520**, 6053–6069 (2023).
- Urry, C. M. & Padovani, P. Unified schemes for radio-loud active galactic nuclei. *Publ. Astron. Soc. Pac.* **107**, 803 (1995).
- McKinney, J. C., Tchekhovskoy, A. & Blandford, R. D. Alignment of magnetized accretion disks and relativistic jets with spinning black holes. *Science* **339**, 49–52 (2013).
- Blandford, R., Meier, D. & Readhead, A. Relativistic jets from active galactic nuclei. *Annu. Rev. Astron. Astrophys.* **57**, 467–509 (2019).
- Jones, D. L., Sramek, R. A. & Terzian, Y. VLBI observations of galactic nuclei. *Astrophys. J.* **246**, 28–37 (1981).
- Mandarakas, N. *Global Alignments of Radio-Jets*. MSc thesis, Univ. of Crete (2019).
- Vazquez Najar, L. A. & Andernach, H. Radio-optical alignment of radio galaxies. Preprint at <https://arxiv.org/abs/1908.09989> (2019).
- Mackay, C. D. The orientation of radio sources and elliptical galaxies. *Mon. Not. R. Astron. Soc.* **151**, 421–425 (1971).
- Palimaka, J. J., Bridle, A. H., Fomalont, E. B. & Brandie, G. W. A preferred orientation for large radio sources relative to their elliptical galaxies. *Astrophys. J. Lett.* **231**, L7–L11 (1979).
- Kinney, A. L. et al. Jet directions in Seyfert galaxies. *Astrophys. J.* **537**, 152–177 (2000).
- Schmitt, H. R., Pringle, J. E., Clarke, C. J. & Kinney, A. L. The orientation of jets relative to dust disks in radio galaxies. *Astrophys. J.* **575**, 150–155 (2002).
- Greenhill, L. J., Kondratko, P. T., Moran, J. M. & Tilak, A. Discovery of candidate H₂O disk masers in active galactic nuclei and estimations of centripetal accelerations. *Astrophys. J.* **707**, 787–799 (2009).
- Lagos, C. D. P., Padilla, N. D., Strauss, M. A., Cora, S. A. & Hao, L. Host galaxy–active galactic nucleus alignments in the Sloan Digital Sky Survey Data Release 7. *Mon. Not. R. Astron. Soc.* **414**, 2148–2162 (2011).
- Battye, R. A. & Browne, I. W. A. Radio and optical orientations of galaxies. *Mon. Not. R. Astron. Soc.* **399**, 1888–1900 (2009).
- Véron-Cetty, M. P. & Véron, P. Are all radio galaxies genuine ellipticals? *Astron. Astrophys.* **375**, 791–796 (2001).

17. Komossa, S. et al. Radio-loud narrow-line type 1 quasars. *Astron. J.* **132**, 531–545 (2006).
18. Janssen, R. M. J., Röttgering, H. J. A., Best, P. N. & Brinchmann, J. The triggering probability of radio-loud AGN. A comparison of high and low excitation radio galaxies in hosts of different colors. *Astron. Astrophys.* **541**, A62 (2012).
19. Wu, Z., Ho, L. C. & Zhuang, M. Y. An elusive population of massive disk galaxies hosting double-lobed radio-loud active galactic nuclei. *Astrophys. J.* **941**, 95 (2022).
20. Kormendy, J. & Ho, L. C. Coevolution (or not) of supermassive black holes and host galaxies. *Annu. Rev. Astron. Astrophys.* **51**, 511–653 (2013).
21. Codis, S. et al. Galaxy orientation with the cosmic web across cosmic time. *Mon. Not. R. Astron. Soc.* **481**, 4753–4774 (2018).
22. Schaye, J. et al. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *Mon. Not. R. Astron. Soc.* **446**, 521–554 (2015).
23. Crain, R. A. et al. The EAGLE simulations of galaxy formation: calibration of subgrid physics and model variations. *Mon. Not. R. Astron. Soc.* **450**, 1937–1961 (2015).
24. Machalski, J., Jamroz, M., Stawarz, Ł. & Weżgowiec, M. Dynamical analysis of the complex radio structure in 3C 293: clues on a rapid jet realignment in X-shaped radio galaxies. *Astron. Astrophys.* **595**, A46 (2016).
25. Barnes, J. E. & Hernquist, L. E. Fueling starburst galaxies with gas-rich mergers. *Astrophys. J. Lett.* **370**, L65 (1991).
26. Plavin, A. V., Kovalev, Y. Y. & Petrov, L. Y. Dissecting the AGN disk–jet system with joint VLBI–Gaia analysis. *Astrophys. J.* **871**, 143 (2019).
27. Kovalev, Y. Y., Petrov, L. & Plavin, A. V. VLBI–Gaia offsets favor parsec-scale jet direction in active galactic nuclei. *Astron. Astrophys.* **598**, L1 (2017).
28. Pashchenko, I. N., Plavin, A. V., Kutkin, A. M. & Kovalev, Y. Y. A bias in VLBI measurements of the core shift effect in AGN jets. *Mon. Not. R. Astron. Soc.* **499**, 4515–4525 (2020).
29. Hopkins, P. F., Hernquist, L., Hayward, C. C. & Narayanan, D. Why are active galactic nuclei and host galaxies misaligned? *Mon. Not. R. Astron. Soc.* **425**, 1121–1128 (2012).
30. Gómez, J. L. et al. Probing the innermost regions of AGN jets and their magnetic fields with RadioAstron. I. Imaging BL Lacertae at 21 microarcsecond resolution. *Astrophys. J.* **817**, 96 (2016).
31. Boccardi, B., Krichbaum, T. P., Bach, U., Bremer, M. & Zensus, J. A. First 3 mm–VLBI imaging of the two-sided jet in Cygnus A. Zooming into the launching region. *Astron. Astrophys.* **588**, L9 (2016).
32. Event Horizon Telescope Collaboration. First M87 Event Horizon Telescope results. I. The shadow of the supermassive black hole. *Astrophys. J. Lett.* **875**, L1 (2019).
33. Lu, R. S. et al. A ring-like accretion structure in M87 connecting its black hole and jet. *Nature* **616**, 686–690 (2023).
34. Janssen, M. et al. Event Horizon Telescope observations of the jet launching and collimation in Centaurus A. *Nat. Astron.* **5**, 1017–1028 (2021).
35. Hodgson, J. A. et al. Location of γ -ray emission and magnetic field strengths in OJ 287. *Astron. Astrophys.* **597**, A80 (2017).
36. Gómez, J. L. et al. Probing the innermost regions of AGN jets and their magnetic fields with RadioAstron. V. Space and ground millimeter–VLBI imaging of OJ 287. *Astrophys. J.* **924**, 122 (2022).
37. Giovannini, G. et al. A wide and collimated radio jet in 3C84 on the scale of a few hundred gravitational radii. *Nat. Astron.* **2**, 472–477 (2018).
38. Oh, J. et al. A persistent double nuclear structure in 3C 84. *Mon. Not. R. Astron. Soc.* **509**, 1024–1035 (2022).
39. Parascos, G. F. et al. Jet kinematics in the transversely stratified jet of 3C 84. A two-decade overview. *Astron. Astrophys.* **665**, A1 (2022).
40. Kim, J. Y. et al. Event Horizon Telescope imaging of the archetypal blazar 3C 279 at an extreme 20 microarcsecond resolution. *Astron. Astrophys.* **640**, A69 (2020).
41. Combes, F. et al. ALMA observations of molecular tori around massive black holes. *Astron. Astrophys.* **623**, A79 (2019).
42. Hodgson, J. A., L’Huillier, B., Liodakis, I., Lee, S.-S. & Shafieloo, A. Using variability and VLBI to measure cosmological distances. *Mon. Not. R. Astron. Soc.* **495**, L27–L31 (2020).
43. Petrov, L. et al. Precise geodesy with the Very Long Baseline Array. *J. Geod.* **83**, 859–876 (2009).
44. Pushkarev, A. B. & Kovalev, Y. Y. Single-epoch VLBI imaging study of bright active galactic nuclei at 2 GHz and 8 GHz. *Astron. Astrophys.* **544**, A34 (2012).
45. Plavin, A. V., Kovalev, Y. Y. & Pushkarev, A. B. Direction of parsec-scale jets for 9220 active galactic nuclei. *Astrophys. J. Suppl. Ser.* **260**, 4 (2022).
46. Malkin, Z. A new version of the OCARS catalog of optical characteristics of astrometric radio sources. *Astrophys. J. Suppl. Ser.* **239**, 20 (2018).
47. Blanton, M. R. et al. Sloan Digital Sky Survey IV: mapping the Milky Way, nearby galaxies, and the distant Universe. *Astron. J.* **154**, 28 (2017).
48. Bundy, K. et al. Overview of the SDSS-IV MaNGA Survey: mapping nearby galaxies at Apache Point Observatory. *Astrophys. J.* **798**, 7 (2015).
49. Abdurro’uf, N. et al. The seventeenth data release of the Sloan Digital Sky Surveys: complete release of MaNGA, MaStar, and APOGEE-2 data. *Astrophys. J. Suppl. Ser.* **259**, 35 (2022).
50. Gunn, J. E. et al. The Sloan Digital Sky Survey photometric camera. *Astron. J.* **116**, 3040–3081 (1998).
51. *Imaging Data Caveats* (Sloan Digital Sky Survey, 2024); [https://www.sdss4.org/dr17/imaging/caveats/#Errorsinmodelmagnitudelocationangles\(phiDevErr,expDevErr\)notreportedinDR8andlater](https://www.sdss4.org/dr17/imaging/caveats/#Errorsinmodelmagnitudelocationangles(phiDevErr,expDevErr)notreportedinDR8andlater)
52. Wolf, C. et al. SkyMapper Southern Survey: first data release (DR1). *Publ. Astron. Soc. Aust.* **35**, e010 (2018).
53. Onken, C. A. et al. SkyMapper Southern Survey: second data release (DR2). *Publ. Astron. Soc. Aust.* **36**, e033 (2019).
54. Source Extractor (SExtractor) v.2.24.2 (IAP/CNRS/UPMC, 2017).
55. Dey, A. et al. Overview of the DESI Legacy Imaging Surveys. *Astron. J.* **157**, 168 (2019).
56. Lang, D. & Hogg, D. W. The Tractor. *GitHub* <https://github.com/dstndstn/tractor> (2020).
57. de Jong, J. T. A. et al. The Kilo-Degree Survey. *The Messenger* **154**, 44–46 (2013).
58. Honscheid, K. & DePoy, D. L. The Dark Energy Camera (DECam). Preprint at <https://arxiv.org/abs/0810.3600> (2008).
59. Flaugher, B. et al. The Dark Energy Camera. *Astron. J.* **150**, 150 (2015).
60. Morganson, E. et al. The Dark Energy Survey image processing pipeline. *Publ. Astron. Soc. Pac.* **130**, 074501 (2018).
61. Abbott, T. M. C. et al. The Dark Energy Survey: Data Release 1. *Astrophys. J. Suppl. Ser.* **239**, 18 (2018).
62. Abbott, T. M. C. et al. The Dark Energy Survey Data Release 2. *Astrophys. J. Suppl. Ser.* **255**, 20 (2021).
63. Sevilla-Noarbe, I. et al. Dark Energy Survey year 3 results: photometric data set for cosmology. *Astrophys. J. Suppl. Ser.* **254**, 24 (2021).
64. Mandarakas, N. et al. Local alignments of parsec-scale AGN radiojets. *Astron. Astrophys.* **653**, A123 (2021).
65. Brainerd, T. G. Anisotropic distribution of SDSS satellite galaxies: planar (not polar) alignment. *Astrophys. J. Lett.* **628**, L101–L104 (2005).

66. L'Huillier, B., Park, C. & Kim, J. Ecology of dark matter haloes – II. Effects of interactions on the alignment of halo pairs. *Mon. Not. R. Astron. Soc.* **466**, 4875–4887 (2017).
67. McAlpine, S. et al. The EAGLE simulations of galaxy formation: public release of halo and galaxy catalogues. *Astron. Comput.* **15**, 72–89 (2016).
68. Thob, A. C. R. et al. The relationship between the morphology and kinematics of galaxies and its dependence on dark matter halo structure in EAGLE. *Mon. Not. R. Astron. Soc.* **485**, 972–987 (2019).
69. de Graaff, A. et al. Observed structural parameters of EAGLE galaxies: reconciling the mass-size relation in simulations with local observations. *Mon. Not. R. Astron. Soc.* **511**, 2544–2564 (2022).
70. Trayford, J. W. et al. Colours and luminosities of $z=0.1$ galaxies in the EAGLE simulation. *Mon. Not. R. Astron. Soc.* **452**, 2879–2896 (2015).
71. Astropy Collaboration. Astropy: a community Python package for astronomy. *Astron. Astrophys.* **558**, A33 (2013).
72. Price-Whelan, A. M. et al. The Astropy project: building an open-science project and status of the v.2.0 core package. *Astron. J.* **156**, 123 (2018).
73. Górski, K. M. et al. HEALPix: a framework for high-resolution discretization and fast analysis of data distributed on the sphere. *Astrophys. J.* **622**, 759–771 (2005).
74. Zonca, A. et al. healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python. *J. Open Source Softw.* **4**, 1298 (2019).
75. Oliphant, T. *Guide to NumPy* (Trelgol Publishing, 2006).
76. Virtanen, P. et al. SciPy 1.0 – fundamental algorithms for scientific computing in Python. Preprint at <https://arxiv.org/abs/1907.10121> (2019).
77. Hunter, J. D. Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95 (2007).

Acknowledgements

We acknowledge useful discussions with D. Blinov, I. Liodakis, I. Sevilla-Noarbe and D. Lang. We also acknowledge the Friends of Sejong 2023 conference (<https://sites.google.com/view/friends-of-sejong2024>), at which most of the results were discussed. D.F.G. and J.A.H. acknowledge the support of the National Research Foundation of Korea (NRF) (NRF-2021R1C1C1009973). B.L. acknowledges the support of the NRF (NRF-2022R1F1A1076338) and the support of the Korea Institute for Advanced Study (KIAS) grant funded by the government of Korea. D.F.G. acknowledges financial support from the Predoctoral Gobierno de Aragón fellowship (2023–2027). J.A. acknowledges the support of the Spanish Ministerio de Universidades and the European Union NextGenerationEU package through the María Zambrano programme (CT33/21) at Universidad Complutense de Madrid (UCM), UCM project PR3/23-30808 and the Ministry of Science, Innovation and Universities (Spain) project PID2022-138263NB-IO (AEI/FEDER, UE). C.S. acknowledges support from the NRF through grant no. 2021R1A2C101302413 funded by the Korean Ministry of Science and ICT. M.J.J. acknowledges support from the NRF through grant nos. 2022R1A2C1003130 and RS-2023-00219959. D.P. acknowledges support from the project ‘Understanding Dark Universe Using Large Scale Structure of the Universe’, funded by the Korean Ministry of Science and ICT. We acknowledge the use of the Astrogeo VLBI FITS image database at http://astrogeo.smce.nasa.gov/vlbi_images, maintained by L. Petrov. The Legacy Surveys consist of three individual and complementary projects: DECaLS (proposal ID no. 2014B-0404; principal investigators: D. Schlegel and A. Dey), BASS (National Optical Astronomy Observatory (NOAO) proposal ID no. 2015A-0801; principal investigators: Z. Xu and X. Fan) and MzLS (proposal ID no. 2016A-0453; principal investigator: A. Dey). DECaLS, BASS and MzLS together include data obtained, respectively, at the

Víctor M. Blanco Telescope, Cerro Tololo Inter-American Observatory, the National Science Foundation’s (NSF’s) National Optical–Infrared Astronomy Research Laboratory (NOIRLab); the Bok Telescope, Steward Observatory, University of Arizona; and the Nicholas U. Mayall Telescope, Kitt Peak National Observatory, NOIRLab. Pipeline processing and analyses of the data were supported by NOIRLab and the Lawrence Berkeley National Laboratory (LBNL). The Legacy Surveys project is honoured to be permitted to conduct astronomical research on Iolkam Du’ag (Kitt Peak), a mountain with particular significance to the Tohono O’odham Nation. NOIRLab is operated by the Association of Universities for Research in Astronomy under a cooperative agreement with the NSF. LBNL is managed by the Regents of the University of California under contract to the US Department of Energy (DOE). This project used data obtained with the DECam, which was constructed by the DES collaboration. Funding for the DES projects has been provided by the US DOE; the US NSF; the Ministry of Science, Innovation and Universities of Spain; the Science and Technology Facilities Council of the UK; the Higher Education Funding Council for England; the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign; the Kavli Institute for Cosmological Physics at the University of Chicago; the Center for Cosmology and AstroParticle Physics at Ohio State University; the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University; Financiadora de Estudos e Projetos; Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro; Conselho Nacional de Desenvolvimento Científico e Tecnológico; the Ministério da Ciência, Tecnologia e Inovação; the Deutsche Forschungsgemeinschaft; and the collaborating institutions in the DES. The collaborating institutions are Argonne National Laboratory; the University of California at Santa Cruz; the University of Cambridge; Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas – Madrid; the University of Chicago; University College London; the DES-Brazil Consortium; the University of Edinburgh; the Eidgenössische Technische Hochschule Zürich; the Fermi National Accelerator Laboratory; the University of Illinois at Urbana-Champaign; the Institut de Ciències de l’Espai (Institut d’Estudis Espacials de Catalunya/ Consejo Superior de Investigaciones Científicas); the Institut de Física d’Altes Energies; LBNL; the Ludwig-Maximilians-Universität München and the associated Excellence Cluster Universe, the University of Michigan; NSF’s NOIRLab; the University of Nottingham; Ohio State University; the University of Pennsylvania; the University of Portsmouth; SLAC National Accelerator Laboratory; Stanford University; the University of Sussex; and Texas A&M University. The BASS is a key project of the Telescope Access Program, which has been funded by the National Astronomical Observatories of China, the Chinese Academy of Sciences (the Strategic Priority Research Program ‘The Emergence of Cosmological Structures’ grant no. XDB09000000) and the Special Fund for Astronomy from the Ministry of Finance. The BASS is also supported by the External Cooperation Program of the Chinese Academy of Sciences (grant no. 114A11KYSB20160057) and the National Natural Science Foundation of China (grant nos. 12120101003 and 11433005). The Legacy Survey team makes use of data products from the Near-Earth Object Wide-field Infrared Survey Explorer (NEOWISE), which is a project of the Jet Propulsion Laboratory/California Institute of Technology. NEOWISE is funded by the National Aeronautics and Space Administration. The Legacy Surveys imaging of the DESI LS footprint is supported by the Director of the Office of Science and the Director of the Office of High Energy Physics of the US DOE under contract no. DE-AC02-05CH1123, by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract, and by the US NSF, Division of Astronomical Sciences, under contract no. AST-0950945 to NOAO. Funding for the SDSS-IV has been provided by the Alfred P. Sloan Foundation, the US DOE

Office of Science and the participating institutions. SDSS-IV acknowledges support and resources from the Center for High Performance Computing at the University of Utah. The SDSS website is www.sdss4.org. SDSS-IV is managed by the Astrophysical Research Consortium for the participating institutions of the SDSS Collaboration, including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, Center for Astrophysics | Harvard & Smithsonian, the Chilean Participation Group, the French Participation Group, Instituto de Astrofísica de Canarias, The Johns Hopkins University, the Kavli Institute for the Physics and Mathematics of the Universe/University of Tokyo, the Korean Participation Group, LBNL, the Leibniz-Institut für Astrophysik Potsdam, the Max-Planck-Institut für Astronomie (Heidelberg), the Max-Planck-Institut für Astrophysik (Garching), the Max-Planck-Institut für Extraterrestrische Physik, the National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatório Nacional/Ministério da Ciência, Tecnologia e Inovação, Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, the UK Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University and Yale University. This project used public archival data from the DES. Funding for the DES projects has been provided by the US DOE; the US NSF; the Ministry of Science, Innovation and Universities of Spain; the Science and Technology Facilities Council of the UK; the Higher Education Funding Council for England; the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign; the Kavli Institute of Cosmological Physics at the University of Chicago; the Center for Cosmology and AstroParticle Physics at Ohio State University; the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University; Financiadora de Estudos e Projetos; Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro; Conselho Nacional de Desenvolvimento Científico e Tecnológico; the Ministério da Ciência, Tecnologia e Inovação; the Deutsche Forschungsgemeinschaft; and the collaborating institutions in the Dark Energy Survey. The collaborating institutions are Argonne National Laboratory; the University of California at Santa Cruz; the University of Cambridge; Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas – Madrid; the University of Chicago; University College London; the DES-Brazil Consortium; the University of Edinburgh; the Eidgenössische Technische Hochschule Zürich; the Fermi National Accelerator Laboratory; the University of Illinois at Urbana-Champaign; the Institut de Ciències de l'Espai (Institut d'Estudis Espacials de Catalunya/Consejo Superior de Investigaciones Científicas); the Institut de Física d'Altes Energies; LBNL; the Ludwig-Maximilians-Universität München and the associated Excellence Cluster Universe; the University of Michigan; the NOAO; the University of Nottingham; Ohio State University; the OzDES Membership Consortium; the University of Pennsylvania; the University of Portsmouth; the SLAC National Accelerator Laboratory; Stanford University; the University of Sussex; and Texas A&M University. DES observations are based in part on observations at the Cerro Tololo Inter-American Observatory, NOAO, which is operated by the Association of Universities for Research in Astronomy under a cooperative agreement with the NSF. The national facility capability for SkyMapper has been funded through Australian Research Council's Linkage Infrastructure, Equipment and Facilities grant (no. LE130100104), awarded to the University of Sydney, the Australian National University, Swinburne University of Technology, the University of Queensland, the University of Western Australia, the University of Melbourne, Curtin University of Technology, Monash University and the Australian Astronomical Observatory. SkyMapper is owned and

operated by the Australian National University's Research School of Astronomy and Astrophysics. The survey data were processed and provided by the SkyMapper Team at the Australian National University. The SkyMapper node of the All-Sky Virtual Observatory is hosted at the National Computational Infrastructure. Development and support of the SkyMapper node of the All-Sky Virtual Observatory has been funded in part by Astronomy Australia Limited and the Australian government through the Commonwealth's Education Investment Fund and National Collaborative Research Infrastructure Strategy, particularly the National eResearch Collaboration Tools and Resources and the Australian National Data Service Projects. We acknowledge the Virgo Consortium for making their simulation data available. The EAGLE simulations were performed using the Distributed Research utilising Advanced Computing-2 facility at Durham, managed by the Institute for Computational Cosmology, and the Partnership for Advanced Computing in Europe facility, CURIE, based in France at the Très Grand Centre de Calcul, the French Alternative Energies and Atomic Energy Commission, Bruyères-le-Châtel. This research made use of Astropy, a community-developed core Python package for astronomy^{71,72}, the HEALPix and healpy package^{73,74}, the NumPy package⁷⁵, the SciPy package⁷⁶ and Matplotlib package⁷⁷.

Author contributions

D.F.G. led and developed the project, produced the code and figures, analysed the data and wrote the paper. J.A.H. designed and led the project and wrote the paper. B.L. helped with the design of the project, carried out the statistical analysis and wrote the paper. J.A. provided help with the optical catalogues and contributed to the data analysis and writing the paper. C.S. prepared and assembled the catalogues of imaging data and the simulation, and co-wrote the sections on them in the paper. K.F. provided guidance to the first author on the intricacies of galaxy shape measurements and the writing of the text. M.J.J. helped with the measurement of galaxy orientation. D.P. helped with the design of the project. F.C. helped to interpret the results and write justification concerning the possible alignment or misalignment between the small-scale radio jets and the optical spin axis of the galaxies.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41550-024-02407-4>.

Correspondence and requests for materials should be addressed to J. A. Hodgson.

Peer review information *Nature Astronomy* thanks Eric Perlman and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2024